Answers for Input-Output

Which of the following statements about the Linux I/O is correct?

A. Incorrect. Applications need to use different virtual file system APIs to access different I/O devices.
B. Correct. Devices can have their own ways to implement file system API.
C. Incorrect. The Linux file system API cannot be used for network communication.

Solution: Linux has a unified virtual file system API for all the I/O devices. But each device can have its specific way of implementing it. This API can not only deals with file accesses on the local machine, but also handles read and write from the network.

If you are an engineer developing a device driver in a high-level language like C, why would you prefer memory-mapped I/O?

A. Incorrect. Memory-mapped I/O is much faster than port-mapped I/O
B. Correct. Port-mapped I/O doesn't support high-level language as it needs special assembly instructions
C. Incorrect. Memory-mapped I/O doesn't need support of memory management unit

Solution: for port-mapped I/O we need to use special assembly instructions, thus we fail to develop it in C. Instead, for memory-mapped I/O, the device register and data buffer can be accessed by general CPU instructions that access specific memory addresses. Hence we can use C to develop the driver. For memory-mapped I/O, we need support from the memory management unit to distinguish I/O requests, which are mapped into virtual address space, from legitimate memory accesses.

If an external device such as a printer runs much slower than the CPU, which mechanism should the driver use to increase efficiency?

A. Incorrect. busy-wait
B. Correct. interrupt

Solution: If the device runs much slower than the CPU, it needs longer time to process the request. If we use busy-wait, most of the time the driver is doing useless check. On the other hand, if we use the interrupt, the driver does not need to actively check all the time. Thus CPU can use this period of time doing something else.

Which of the following statements about the DMA is incorrect?

A. Incorrect. DMA is a hardware core inside CPU
B. Correct. CPU is not involved in the data transfer until it receives a hardware interrupt
C. Correct. Multiple device controllers can share the same DMA

Solution. The reason we use DMA is to allow an input/output (I/O) device to send or receive data directly to or from the main memory, bypassing the CPU. Thus, it couldn't be one part of the CPU. The CPU is involved only when the DMA generates a hardware interrupt. As a DMA is so useful these days, there could be one on the parent motherboard for all devices.

Which of the following statement(s) about device drivers are correct? Select all that apply.

A. Incorrect. Device driver is implemented in the firmware of the external device

B. Correct. Device drivers can come prepackaged with the OS or be installed dynamically.

C. Correct. Device driver starts I/O operation and blocks until the interrupt handler unblocks it when I/O completes.

D. Correct. Device driver unblocks the user process that calls the I/O system call when I/O completes.

Solution. A device driver is implemented in the OS. On the other hand, the device controller is a piece of the external device.

Which of the following statement(s) about devices is(are) true? Select all that apply.

A. Correct. With the help of Linux file system API, we can map every device into a file name, and access it as a file on the disk.

B. Correct. User processes can directly configure devices using ioctl system call.

C. Incorrect. The scheduling of I/O devices is the same as the scheduling of the CPU.

Solution: On Linux, we can name each device as a file, and access it by Linux file system API. I/O control(ioctl) system call is a way for user processes to directly control the external devices. Each device has its own I/O scheduler. And these schedulers are different from the one used for CPU scheduling. An I/O scheduler decides which process's I/O operation request gets sent to the external device, unlike the CPU scheduler that decides which running process gets the next CPU cycle.

Which of the following statements about the disk partition is correct?

A. Incorrect. Each partition contains its own partition table.

B. Correct. If we want to run both Ubuntu and Windows on our laptop, we need to break up the disk drive into multiple partitions.

C. Incorrect. Each partition has a boot block, from which we can get information about the entire drive.

Solution: A partition table at a fixed location on disk tells us the beginning of each partition. There is only one table for all partitions.
Typically, disk drives can be broken up into several logical drives called partitions. Multiple partitions allow us to run different operating systems in each partition. The master boot record is stored at the beginning of the partition table, which stores information about the entire drive.

What is the advantage of linked list allocation compared to contiguous allocation?

A. Correct. Linked list allocation doesn't cause external fragmentation.

B. Incorrect. Linked list allocation provides faster random access.

C. Incorrect. Linked list allocation makes sure that the content of one file is stored in sequential physical blocks.

Solution: Linked list allocation doesn't cause external fragmentation, but it is very slow. If one has a file that consists of multiple blocks and all I'm interested in is to read the nth block, this technique will require first reading the initial n-1 blocks up front. Note that logically each file has contiguous blocks, but physically on the hard disk, the file is spread out into different physical blocks.

Which of the following statements about the FAT technique is correct? Select all possible answers.

A. Correct. FAT is loaded into the main memory whenever we boot up the operating system.
B. Correct. FAT is a technique to address the issue that linked list allocation is slow for random access.
C. Correct. One of the drawbacks of FAT is that when the disk gets large, the FAT table may be too big to fit in memory.

Solution: FAT is much faster for random access compared to linked list allocation, because it's loaded into the main memory, and the traversing time in the main memory is much faster than on the disk. But the problem with FAT is when disk size increases, the size of the FAT table also grows. If the disk is too big, the FAT table will fail to fit in the memory.

Which of the following statements about i-nodes is correct? Select all possible answers.

A. Correct. The memory requirement is far less than FAT, because we only need to load i-nodes of open files in memory.
B. Incorrect. For i-nodes design, both small and big file accesses are fast.
C. Correct. We need more disk I/O if we want to fetch content at the end of a big file, which uses triple indirect block.

Solution:
The memory requirement of i-nodes is far less than FAT because we only need to load i-nodes of open files in memory. In the i-node design, small file access is fast, but big file access can take a longer time due to use the indirect blocks. This also means that the larger the file is, the more disk I/O that you would need to fetch.

Which of the following statements about directories is correct?

A. Correct. A directory entry contains information such as the file name and the address of the first block of the file.
B. Incorrect. Directories are stored as special blocks in the file system.
C. Incorrect. We can use the in-line approach to store long file names on the heap.

Solution:
A directory entry provides the information needed to find the disk data blocks of a file. For DOS and Unix, each entry contains a pair of the file name and the address of the first block of the file. Directories are stored as regular blocks in the file system. We can use the heap approach to deal with long file names, which just stores a pointer to the name in the directory entry, and the actual file name is stored in the heap.

Which of the following statements about links is correct? Select all possible answers.

A. Correct. Links are used for sharing files between different users or creating shortcuts.
B. Incorrect. In hard links, we maintain two separate i-nodes if there are two different paths to the same file.
C. Correct. In soft links, we can delete the i-node for the actual file even if there is another i-node of type LINK containing its path.

Solution:
Links are used for sharing files between different users or creating shortcuts. There are two ways to do it: one is hard links and the other is soft links. In hard links, if there are two different paths for a file, we create one single i-node and two directory entries that can point to the same i-node. Soft links, on the other hand, maintains two separate i-

nodes. The one that's linking to the original file will instead have an i-node of type LINK, which just contains the path name. As these i-nodes are independent, we can delete anyone without influencing the other.

Which of the following statements about the caching is correct? Select all possible answers.

A. Correct. Caching reduces the chance that the operating system is blocked by the disk I/O.
B. Correct. Caching would cause an inconsistent state between content in the main memory and on the disk.
C. Incorrect. Caching with write though will reduce the performance of both read and write.

Solution:
With caching we don't need to access the disk for every read and write, which makes both processes much faster. But in this case, the blocks that are cached in memory may not be reflected on disk. And hence we will have an inconsistent state. Thus, if we modify important data, we need to use write-through technique to synchronize the state on the disk with the one in the main memory, which will reduce the speed of writing into the file but is still able to optimize read.

Which of the following statements about the logical dump is correct? Select all possible answers.

A. Incorrect. In a logical dump, we copy physically block by block from one disk to another.
B. Correct. In a logical dump, we can copy files between different operating systems.
C. Correct. In a logical dump, we need to copy all the modified directories as well as the files themselves.

Solution:
In a physical backup, we would actually copy physically block by block from one disk to another. Logical backup or logical dump does it more at a file system level. It is agnostic to the operating system. For example, when you're synchronizing to Dropbox, they're actually copying not at the block level, but at the file level, one file at a time.

Based on the blocks in use list and free blocks list below, what kind of consistency issue does the file system have?

Index	0	1	2	3	4	5	6
Blocks in use	1	1	1	0	1	0	1
Free blocks	0	0	0	1	0	0	0

A. Correct. Missing block
B. Incorrect. Duplicate block in free list
C. Incorrect. Duplicate data block

Solution: A

Feedback: block number 5 is neither in use nor free. It means block number 5 is missed by the file system.

Is the following statement true or false?

The CD-ROM uses contiguous allocation scheme for storing files.

A. Correct. True
B. Incorrect. False

Solution:
Contiguous allocation means each file occupies a contiguous region of blocks. It is mostly used in cases where we write once read multiple times, making the content immutable. CD-ROMs, DVD-ROMs are examples of contiguous allocation. Remember that ROM means Read-Only Memory.

Is the following statement true or false?

The linked-list file allocation scheme results in no external fragmentation.
A. Correct. True
B. Incorrect. False

Solution: The statement is true because external fragmentation is created when a file system needs to use a larger amount of disk space to accommodate a smaller file, for reasons such as poor handling of files that have shrunk in size. Linked list file allocation scheme does not incur external fragmentation because we can simply use pointers to point to the empty space blocks when we want to allocate new space.

Is the following statement true or false?

A logical dump always dumps all directories even if they are unmodified.
A. Incorrect. True
B. Correct. False
Solution: Logical dumps do not dump all directories. The files that need backup are the ones that have been modified since the last backup, and the path of directories from the directory that contains the modified file to the root are the directories that have been affected by the modification. Thus, their i-nodes and blocks need to be backed up too. But directories whose files have not changed do not need to have their i-nodes and directory blocks backed up.

Is the following statement true or false?

In UNIX, cached i-node blocks are updated using write-through caching techniques.
A. Correct. True
B. Incorrect. False
Solution: Write-through caching means that once a block has been read from disk into caches in memory, we cannot write to it but can only read it. Once a write operation appears, we must flush the block onto disk in order to maintain file system consistent with in-memory cache. The goal is to maintain consistency on only crucial disk blocks to keep the filesystem in sync, while allowing some cache disparity boost the speed of a write operation. Therefore, in UNIX, write-through caching is used only for I-node blocks, indirect address blocks, free space management blocks, and data blocks are allowed to not have write-through caching.

Is the following statement true or false?

DMA provides direct access between memory and an I/O device.
A. Correct. True
B. Incorrect. False

Solution: Implied by the name, DMA enables the CPU to initiate a data transfer process with the memory address of the data to be written (or address of memory to fill in once data from device is read), what data to read from the device, content length, etc. Then the CPU is free, and the actual data transfer is handled by the DMA channel and the device. When the transfer is complete the CPU gets an interrupt by the DMA controller, so the CPU knows where to find the data in memory. This implies a mapping between memory and data on the device.

Is the following statement true or false?

Directory entries are stored in inodes.
A. Incorrect. True
B. Correct. False

Solution: Directory entries are like file contents except they are understood by the file system. Therefore the entries are stored in file blocks, so the directory file has its own inode as well. But the directory data is not stored in inodes.

Is the following statement true or false?

The boot block is used to load the OS into kernel space.
A. Correct. True
B. Incorrect. False

Solution: The boot block is useful to load the OS code into kernel memory to be run on the CPU. Therefore it is usually stored in a special partition on disk for stability reasons.